

### Telephonic Interview on July 23, 2007

Applicant appreciates the time that Examiner took to conduct a telephonic interview with the undersigned on July 23, 2007. Also present during the telephonic interview was Karl Ginter, who is a technical consultant of the assignee. During the interview, the following two related applications were discussed -- 10/086,602 and 10/086,268.

No exhibits were shown nor were any demonstrations conducted. No particular claims were discussed. However, the feature of an "enterprise directory" in the claims, including the enterprise directory being "a directory of named objects, including users, network devices and network services" was discussed.

The specific prior art that was discussed includes U.S. Patent 6,909,708 to Krishnaswamy et al. and assigned (on its face) to MCI Communications Corporation.

No proposed amendments were discussed.

The principal argument of the Applicant is that the "database" disclosed by the Krishnaswamy patent does not anticipate the "enterprise directory" feature recited in the claims. The Examiner argued that Applicant's specification does not distinguish between a "database" and an "enterprise directory" and that, therefore, the Examiner is free to apply the "database" of Krishnaswamy against the "enterprise directory" recited in the claims.

No agreement was reached during the interview.

### REMARKS

#### Statutory Subject Matter

The Examiner contends that claim 20 is directed to a data structure of a database that does not fall within any of the four categories of statutory subject matter. Applicant respectfully traverses the rejection.

In particular, claim 20 is directed to an article of manufacture. Furthermore, the "enterprise directory" (not a database, as discussed below with respect to the prior art based rejections) embodied thereon produces functional advantages that disappear if the same data is recorded in a different format. According to the In Re Lowry case and cases following it, this is sufficient to distinguish a claimed invention from what might otherwise be considered unpatentable "printed matter." More particularly, that the enterprise directory embodied on the article of manufacture produces functional advantages that disappear if the same data is recorded

in a different format indicates that the claimed format has a functional significance that can be separated from its data recording contents.

Examples of such functional significance pervade Applicant's specification. For example, see paragraph [0086] of the published version of Applicant's specification, which discusses:

[0086] In a preferred embodiment, the enterprise directory 90 is implemented using a general purpose directory such as NDS. The invention introduces schema extensions in NDS. The schema extensions enhance the NDS base schema so that it supports the Directory Services requirements for an H.323 Recommendation based IP telephony network. In addition to H.323 support, the schema extension enables the H.323 gatekeepers in an H.323 IP telephony network to automatically find each other. This capability is not currently specified and supported by the ITU H.323 Recommendation v.1. The schema extensions also enable the invention to provide many of its unique features, e.g. caller ID, follow me with call filtering, etc.

Since the claimed format has a functional significance that can be separated from its data recording contents, claim 20 recites patentable subject matter under 35 USC 101.

#### Indefiniteness

Claims 28 and 38 are rejected as being indefinite, due to the use of a trademark. Claim 28 has been cancelled. Claim 38 had been amended to remove the portion reciting the trademarks.

#### Prior Art Based Rejections

It is noted that the Krishnaswamy patent is either the sole reference or a primary reference in the prior art based rejections. In using the Krishnaswamy patent to reject the claims, the Examiner is contending that Krishnaswamy's disclosure of a "database" anticipates the "enterprise directory" feature in Applicant's claims. Note that the Examiner does not contend that Krishnaswamy's disclosure of a database renders obvious the "enterprise directory" feature in Applicant's claims. For example, in rejecting claims 21, 27 and 29-31 of the '602 application as anticipated under 35 U.S.C. 102(e)<sup>1</sup>, the Examiner notes in part (see page 3 of Office Action, in Item 6):

---

<sup>1</sup> It is noted that 35 USC 102(e) is not applicable in any event. Applicant is treating this rejection as if a rejection under 35 USC 102(a).

... and an enterprise directory server (Fig. 19F, Directory is enterprise directory server or Fig. 10A, Ref 1, 2, 3) coupled to the plurality of gateway networks, the enterprise directory server comprising an enterprise directory that is a directory of named objects, including users, network devices and network services and having an extensible schema configured to provide data to support routing of telephone calls (Fig. 10B, Ref enterprise Directory server 1082, user profile, telephone gateway and email address etc. and Fig. 10A, Ref 1-3).

As mentioned above in the Examiner Interview Summary, the Examiner has argued that Applicant's specification does not distinguish between a "database" and an "enterprise directory" and that, therefore, the Examiner is free to apply the "database" of Krishnaswamy against the "enterprise directory" recited in the claims.

It is respectfully submitted that there is no such requirement for Applicant's specification to distinguish between a "database" and an "enterprise directory." If this was the case, then applicants would be required to predict every rejection an Examiner may present and provide rebuttal distinguishing language in the specification, even before the patent application has been filed. Clearly, this cannot be the requirement.

Thus, Applicant once again respectfully submits that the "database" of Krishnaswamy does not anticipate the "enterprise directory" recited in the claims. It is respectfully submitted that it is the Examiner who has the initial burden to demonstrate that each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference. Here, the Examiner has not demonstrated that the Krishnaswamy "database" anticipates the "enterprise directory" recited in the claims. It is respectfully submitted that the anticipation rejection is improper for at least this reason alone.

However, even though it is not Applicant's burden to do so, in the interest of furthering prosecution and supplementing the record, Applicant provides herewith evidence that the Krishnaswamy "database" does not anticipate the "enterprise directory" recited in the claims. Applicant does this by way of providing several references from about the time of Applicant's priority date.

1. T. Howes and M. Smith, LDAP: Programming Directory Enabled Applications with Lightweight Directory Access Protocol. Macmillan Technology Series, 1997. In addition to the cover and publication pages, Applicant provides herewith an excerpt at pages 4-5. See Appendix A to this Amendment C. In the excerpt, in the section entitled "What a Directory Service is Not," it is clearly stated that a directory services is not a general purpose database. While this

entire section is provided in Appendix A, Applicant would like to draw the Examiner's attention to one particular paragraph:

The first thing you should not do is treat the directory like a general database. It's not designed with that kind of use in mind, and you'll likely be disappointed with the results. This means the directory is not suited to provide transaction-based service, cannot generally handle huge numbers of updates, and usually would make a bad engine for an SQL application.

2. United States Patent No. 6,016,499, assigned on its face to Novell, Inc. The filing date of the priority provisional application is February 20, 1997, and the filing date of the non-provisional application is July 21, 1997. At col. 1, line 60 et seq, for example, the applicant distinguishes between "directory services" and "relational databases."

Unlike relational databases, "directory service" structures are a relatively recent development in information technology. The need for directory services became clear only after computer networks became a prominent part of the information landscape and grew so large and complex that their administration required full-time attention from specialists. Directory services are sometimes referred to as "naming services."

A variety of directory service providers are now available to help administer both the location of network resources and the rights of network users to use those resources. Many, but not all, directory service tools conform at least in part with the X.500 directory services standard. One well-known directory service system includes NetWare Directory Services software commercially available from Novell, Inc. of Provo, Utah (NETWARE DIRECTORY SERVICES is a trademark of Novell, Inc.). As used herein, "Novell Directory Services" ("NDS") includes NetWare Directory Services and any other directory service commercially available from Novell, Inc.

In contexts other than the present one, a directory services repository is sometimes called a directory services "database." Both repositories and databases may be distributed, because that is mainly a matter of storage and locks for enforcing consistency, rather than a question of the basic internal structure. But "database" and "repository" are not interchangeable in the present context.

A repository supports naming services. Each repository is organized hierarchically, as one or more trees. Each object in a tree (except the root object) has exactly one parent. Objects may generally have children, which may inherit properties from their ancestor objects. Objects are instances of "classes," as described in detail below.

By contrast, "database" as used herein refers to a relational database. A given database may support any of a wide variety of commercial or personal activities. Each database is organized as a set of tables in which rows represent records and columns represent record fields. Certain fields may be found in multiple tables, and the values of these "key" or "index" fields are used to guide database searches. Database access and manipulation involve combining information from various tables into new combinations to obtain different views of the data.

In short, databases and repositories arose at different times to meet different needs, and have different structures and capabilities. There is no need to dwell further here on the details of relational database structure, SQL, or ODBC, as numerous references on those topics are widely available. Many computer science professionals have also taken at least one course in databases.

3. The SLAPD and SLURPD Administrator's Guide, University of Michigan, 30 April 1996, Release 3.3. In addition to the cover, publication pages and Table of Contents, Applicant provides herewith an excerpt at pages 1-10. See Appendix B to this Amendment C. In the excerpt, in the section entitled "1.1 What is a directory service?" it is discussed how a directory service differs from a database. The entire section 1.1 (two paragraphs only), is reproduced below:

**1.1 What is a directory service?**

A directory is like a database, but tends to contain more descriptive, attribute-based information. The information in a directory is generally read much more often than it is written. As a consequence, directories don't usually implement the complicated transaction or roll-back schemes regular databases use for doing high-volume complex updates. Directory updates are typically simple all-or-nothing changes, if they are allowed at all. Directories are tuned to give quick-response to high-volume lookup or search operations. They may have the ability to replicate information widely in order to increase availability and reliability, while reducing response time. When directory information is replicated, temporary inconsistencies between the replicas may be OK, as long as they get in sync eventually.

There are many different ways to provide a directory service. Different methods allow different kinds of information to be stored in the directory, place different requirements on how that information can be referenced, queried and updated, how it is protected from unauthorized access, etc. Some directory services are local, providing service to a restricted context (e.g., the finger service on a single machine). Other services are global, providing service to a much broader context (e.g., the entire Internet). Global services are usually distributed, meaning that the data they contain is spread across many machines, all of which cooperate to provide the directory service. Typically a global service defines a uniform namespace which gives the same view of the data no matter where you are in relation to the data itself.

Applicant has thus amply demonstrated that the Krishnaswamy "database" does not anticipate the "enterprise directory" recited in the claims. Since the Examiner's contention regarding the Examiner's contention regarding the Krishnaswamy "database" anticipating the "enterprise directory" pervades all the prior art based rejections, it is respectfully submitted that all of the prior art based rejections should be withdrawn.

## CONCLUSION

Applicants' believe that all pending claims are allowable and respectfully request a Notice of Allowance for this application from the Examiner. Should the Examiner believe that a telephone conference would expedite the prosecution of this application, the undersigned can be reached at the telephone number set out below.

Respectfully submitted,

BEYER WEAVER LLP

/ASH/

Alan S. Hodes


Reg. No. 38,185

P.O. Box 70250  
Oakland, CA 94612-0250  
408-255-8001

## APPENDIX A



M  
TP



MACMILLAN TECHNOLOGY SERIES

# LDAP

*Programming Directory-Enabled  
Applications with Lightweight  
Directory Access Protocol*

Timothy A. Horvath, Ph.D.  
Mark C. Smith



# **LDAP:**

## Programming Directory-Enabled Applications with Lightweight Directory Access Protocol

Timothy A. Howes

Mark C. Smith



MACMILLAN  
TECHNICAL  
PUBLISHING

Macmillan Technical Publishing, Indianapolis, Indiana



# **LDAP: Programming Directory-Enabled Applications with Lightweight Directory Access Protocol**

By Timothy A. Howes and Mark L. Smith

Published by  
Macmillan Technical Publishing  
201 West 103rd Street  
Indianapolis, IN 46290 USA

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

Copyright © 1997 by Macmillan Technical Publishing

Printed in the United States of America 4 5 6 7 8 9 0

Library of Congress Catalog Card Number: 96-78810

ISBN: 1-57870-000-0

## **Warning and Disclaimer**

This book is designed to provide information about the LDAP computer program. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an "as is" basis. The author(s) and Macmillan Technical Publishing shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the disks or programs that may accompany it.

*Publisher* Don Fowler  
*Publisher's Assistant* Rosemary Lewis  
*Publishing Manager* Jim LeValley  
*Marketing Manager* Kearney Simpson  
*Managing Editor* Carla Hall

**Acquisitions Editor**  
Jim LeValley

**Senior Editors**  
Sarah Kearns  
Juzanne Snyder

**Development Editor**  
Tim Haddleton

**Project/Copy Editor**  
Mizzi Foster

**Acquisitions Coordinator**  
Amy Lewis

**Cover Designer**  
Sandra Schroeder

**Cover Production**  
Aren Howell

**Book Designer**  
Sandra Schroeder

**Manufacturing Coordinator**  
Brook Farling

**Production Manager**  
Kelly Dobbs

**Production Team Supervisors**  
Laurie Casey, Joe Millay

**Graphics/Image Specialists**  
Dan Harris  
Laura Robbins  
Dennis Sheehan

**Production Analyst**  
Erich J. Richter

**Production Team**  
Tina Brown, Teresa Clodder,  
Christopher Morris, Megan Wade

**Indexer**  
Eric Brinkman



While directories come in all shapes and sizes, most examples share some common elements that bind them together. Directories are special purpose databases, usually containing typed information. Read access to the directory is typically much more frequent than update access. Some directories do not allow updates at all, although LDAP does. Directory services are distinguished from name services by their ability to search for, as well as retrieve, named information. They are distinguished from general databases by their higher access/update ratio, their lack of transaction semantics, and their limited ability to separate search from retrieval. The concept of "result set," so useful in a general database system, is not often found in a directory service.

Directories may be replicated for improved reliability and performance, with either strong or weak consistency requirements. Directories may include authentication and access control capabilities, or they may only provide world-readable information that does not require such protection.

Many directories, including LDAP-based ones, essentially provide a way to name, manage, and access collections of attribute-value pairs. Where directories differ from one another is in the way in which this information is represented and accessed, the flexibility with which the information can be searched, whether the kinds of information in the directory may be extended, and whether and how the information can be updated.

The preceding paragraphs provide a description of a directory service, not a formal definition. It is intentionally broad, though we will narrow our definition later when we focus in on LDAP directory service, the subject of this book.

## What a Directory Service Is Not

Having a directory service capable of storing arbitrary attribute-value pairs is an attractive thing. So attractive, in fact, that you may be tempted to put some things in it that you should not, or to use the directory for purposes for which it is not well suited. So the question naturally arises, "What should you *not* do with a directory service?"

Dir



The first thing you should not do is treat the directory like a general database. It's not designed with that kind of use in mind, and you'll likely be disappointed with the results. This means the directory is not suited to provide transaction-based service, cannot generally handle huge numbers of updates, and usually would make a bad engine for an SQL application.

Second, a directory is not a file system. Storing all the files on your system in the directory is probably a bad idea. In fact, any time you are thinking of storing very large objects in the directory, you should ask yourself whether it might not be more appropriate to store only a pointer to the object in the directory, and the object itself in a service more suited to storing and retrieving large objects (for example, on an FTP server or in a file system).

Finally, a directory should not be used unless it's providing your application some benefit. Does the information you want to store in the directory need to be accessed by others? By you or your application from different locations? If it is only to be accessed by one application on a specific host, then there is little benefit derived from storing it in the directory, yet there could be overhead, added complexity, and other factors that argue for a local solution.

## Directory-Enabled Applications

For a few applications—those that read, write, or manage directory information as their primary purpose—the benefit of directory enabling is clear. The applications could not exist without it, after all. But for the vast majority of applications, directory support is ancillary to the application's main purpose, which might be browsing the World Wide Web; sending, receiving, and processing e-mail; organizing information on your intranet; or just about anything else you can think of. It is this latter category of applications that represents the vast majority of potential directory client applications that might benefit from being *directory-enabled*.

## APPENDIX B

# **The SLAPD and SLURPD Administrator's Guide**

University of Michigan

30 April 1996  
Release 3.3



## **Copyright**

Copyright © 1992-1996 Regents of the University of Michigan. All Rights Reserved.

Redistribution and use in source and binary forms are permitted provided that this notice is preserved and that due credit is given to the University of Michigan at Ann Arbor. The name of the University may not be used to endorse or promote products derived from this software or documentation without specific prior written permission. This software is provided "as is" without any express or implied warranty.

## **Acknowledgments**

The LDAP development team at the University of Michigan consists of Tim Howes, Mark Smith, Gordon Good, Lance Sloan and Steve Rothwell. Our thanks also to Bryan Beecher, Frank Richter, Eric Rosenquist, Peter Whittaker, Martijn Koster, Craig Watkins, Rocky Rakesh Patel, Alan Young, Mark Prior, Enrique Silvestre Mora, Roland Hedberg, and numerous others.

# Table of Contents

<b>1. INTRODUCTION TO SLAPD AND SLURPD .....</b>	<b>6</b>
1.1 WHAT IS A DIRECTORY SERVICE?.....	6
1.2 WHAT IS LDAP?.....	6
1.3 HOW DOES LDAP WORK?.....	8
1.4 WHAT IS SLAPD AND WHAT CAN IT DO?.....	8
1.5 WHAT ABOUT X.500?.....	9
1.6 WHAT IS SLURPD AND WHAT CAN IT DO?.....	9
<b>2. A QUICK-START GUIDE TO RUNNING SLAPD.....</b>	<b>10</b>
<b>3. THE BIG PICTURE - CONFIGURATION CHOICES.....</b>	<b>12</b>
3.1 LDAP AS A LOCAL SERVICE ONLY .....	12
3.2 LOCAL SERVICE WITH X.500 REFERRALS.....	12
3.3 LDAP AS A FRONT END TO X.500.....	13
3.4 REPLICATED SLAPD SERVICE .....	13
<b>4. BUILDING AND INSTALLING SLAPD &amp; SLURPD.....</b>	<b>14</b>
4.1 PRE-BUILD CONFIGURATION.....	14
4.1.1 Editing the Make-common file.....	14
4.1.2 Editing the include/ldapconfig.h file.....	16
4.2 MAKING THE SOFTWARE.....	17
4.3 INSTALLING THE SOFTWARE .....	17
<b>5. THE SLAPD CONFIGURATION FILE.....</b>	<b>19</b>
5.1 CONFIGURATION FILE FORMAT.....	19
5.2 CONFIGURATION FILE OPTIONS.....	19
5.2.1 Global Options.....	20
5.2.2 General Backend Options .....	22
5.2.3 LDBM Backend-Specific Options .....	24
5.2.4 Shell Backend-Specific Options.....	25
5.2.5 Password Backend-Specific Options .....	26
5.3 ACCESS CONTROL.....	26
5.3.1 What to control access to.....	26
5.3.2 Who to grant access to .....	27
5.3.3 The access to grant .....	28
5.3.4 Access Control Evaluation.....	28
5.3.5 Access Control Examples .....	28
5.4 SCHEMA ENFORCEMENT.....	29
5.5 CONFIGURATION FILE EXAMPLE .....	30
<b>6. RUNNING SLAPD.....</b>	<b>33</b>
6.1 COMMAND-LINE OPTIONS.....	33
6.2 RUNNING SLAPD AS A STAND-ALONE DAEMON.....	34
6.3 RUNNING SLAPD FROM INETD .....	34
<b>7. MONITORING SLAPD.....</b>	<b>35</b>
<b>8. DATABASE CREATION AND MAINTENANCE TOOLS.....</b>	<b>37</b>
8.1 CREATING A DATABASE OVER LDAP.....	37
8.2 CREATING A DATABASE OFF-LINE.....	38
8.2.1 The ldif2ldbm program.....	39
8.2.2 The ldif2index program.....	40
8.2.3 The ldif2id2entry program.....	41
8.2.4 The ldif2id2children program .....	41

8.2.5 The <i>ldbmcat</i> program .....	41
8.2.6 The <i>ldif</i> program .....	41
8.3 THE LDIF TEXT ENTRY FORMAT .....	42
8.4 CONVERTING FROM QUIPU EDB FORMAT TO LDIF FORMAT .....	43
8.4.1 The <i>edb2ldif</i> program .....	43
8.4.2 Step-by-step EDB to LDIF conversion .....	44
8.5 THE LDBMTEST PROGRAM .....	45
8.6 THE LDBM DATABASE FORMAT .....	46
8.6.1 Overview .....	46
8.6.2 Attribute index format .....	47
8.6.3 Other indexes .....	47
<b>9. PERFORMANCE TUNING.....</b>	<b>48</b>
9.1 THE ALLIDS THRESHOLD .....	48
9.2 THE ENTRY CACHE .....	48
9.3 THE DB CACHE .....	48
9.4 MAINTAIN THE RIGHT INDICES .....	49
<b>10. DISTRIBUTING SLAPD DATA.....</b>	<b>50</b>
<b>11. REPLICATION WITH SLURPD.....</b>	<b>51</b>
11.1 OVERVIEW .....	51
11.2 REPLICATION LOGS .....	51
11.3 COMMAND-LINE OPTIONS .....	52
11.4 CONFIGURING SLURPD AND A SLAVE SLAPD INSTANCE.....	53
11.4.1 Set up the master slapd.....	53
11.4.2 Set up the slave slapd.....	54
11.4.3 Shut down the master slapd .....	54
11.4.4 Copy the master slapd's database to the slave .....	54
11.4.5 Configure the master slapd for replication.....	54
11.4.6 Restart the master slapd and start the slave slapd .....	55
11.4.7 Start slurpd.....	55
11.5 ADVANCED SLURPD OPERATION .....	55
11.5.1 Replication errors .....	55
11.5.2 Slurpd's one-shot mode and reject files.....	56
11.6 REPLICATION FROM A SLAPD DIRECTORY SERVER TO AN X.500 DSA .....	56
<b>12. APPENDIX A: WRITING A SLAPD BACKEND .....</b>	<b>58</b>
12.1 THE SLAPD BACKEND API.....	59
12.1.1 Bind .....	59
12.1.2 Unbind .....	60
12.1.3 Compare .....	60
12.1.4 Search .....	61
12.1.5 Modify .....	63
12.1.6 Modify RDN .....	64
12.1.7 Add .....	65
12.1.8 Delete .....	65
12.1.9 Abandon .....	66
12.1.10 Initialization .....	66
12.1.11 Configuration.....	66
12.1.12 Close.....	67
12.2 UTILITY ROUTINES YOUR BACKEND MAY WANT TO CALL .....	67
12.2.1 Sending Search Entries .....	67
12.2.2 Sending a Result .....	68
12.2.3 Testing a Filter Against an Entry.....	68
12.2.4 Creating an Entry .....	68
<b>13. APPENDIX B: WRITING A SHELL BACKEND.....</b>	<b>70</b>

13.1 OVERVIEW .....	70
13.2 INPUT FORMAT .....	70
13.2.1 Bind.....	70
13.2.2 Unbind.....	71
13.2.3 Search.....	71
13.2.4 Compare.....	71
13.2.5 Modify.....	72
13.2.6 Modify RDN.....	72
13.2.7 Add.....	72
13.2.8 Delete.....	72
13.2.9 Abandon.....	73
13.3 OUTPUT FORMAT .....	73
13.3.1 Search Entry.....	73
13.3.2 Result .....	73
13.3.3 Debugging .....	73
13.4 EXIT STATUS .....	73
13.5 EXAMPLE.....	74
13.5.1 Configuration file.....	74
13.5.2 Search command shell script.....	74
<b>14. APPENDIX C: DISTRIBUTED INDEXING WITH <i>CENTIPEDE</i>.....</b>	<b>76</b>
14.1 AN EXAMPLE.....	77
14.2 LIMITATIONS.....	78
<b>15. APPENDIX D: USING KERBEROS AUTHENTICATION WITH <i>SLAPD</i> AND <i>SLURPD</i>. .....</b>	<b>79</b>
15.1 BUILD THE U-M LDAP PACKAGE WITH KERBEROS SUPPORT ENABLED.....	79
15.2 USING KERBEROS WITH SLAPD.....	79
15.2.1 Obtain a <i>srvtab</i> File for Your <i>slapd</i> Server.....	79
15.2.2 Install the <i>srvtab</i> File and Tell <i>slapd</i> Where It Is .....	80
15.2.3 Add Kerberos Names to Entries to Enable Authentication.....	80
15.2.4 Associate a Kerberos Name with the “ <i>rootdn</i> ” (optional).....	81
15.3 USING KERBEROS WITH SLURPD.....	81
15.3.1 Obtain a <i>srvtab</i> File for Your <i>slurpd</i> Server.....	81
15.3.2 Configure the <i>slapd</i> Slaves to Accept Kerberos Authentication .....	81
15.3.3 Configure <i>slurpd</i> to Use Kerberos When Connecting to the Slaves.....	82

## 1. Introduction to *slapd* and *slurpd*

This document describes how to build, configure, and run the stand-alone LDAP daemon (*slapd*) and the stand-alone LDAP update replication daemon (*slurpd*). It is intended for newcomers and experienced administrators alike. This section provides a basic introduction to directory service, and the directory service provided by *slapd* in particular.

### 1.1 What is a directory service?

A directory is like a database, but tends to contain more descriptive, attribute-based information. The information in a directory is generally read much more often than it is written. As a consequence, directories don't usually implement the complicated transaction or roll-back schemes regular databases use for doing high-volume complex updates. Directory updates are typically simple all-or-nothing changes, if they are allowed at all. Directories are tuned to give quick-response to high-volume lookup or search operations. They may have the ability to replicate information widely in order to increase availability and reliability, while reducing response time. When directory information is replicated, temporary inconsistencies between the replicas may be OK, as long as they get in sync eventually.

There are many different ways to provide a directory service. Different methods allow different kinds of information to be stored in the directory, place different requirements on how that information can be referenced, queried and updated, how it is protected from unauthorized access, etc. Some directory services are *local*, providing service to a restricted context (e.g., the finger service on a single machine). Other services are global, providing service to a much broader context (e.g., the entire Internet). Global services are usually *distributed*, meaning that the data they contain is spread across many machines, all of which cooperate to provide the directory service. Typically a global service defines a uniform *namespace* which gives the same view of the data no matter where you are in relation to the data itself.

### 1.2 What is LDAP?

*Slapd*'s model for directory service is based on a global directory model called LDAP, which stands for the Lightweight Directory Access Protocol. LDAP is a directory service protocol that runs over TCP/IP. The nitty-gritty details of LDAP are defined in RFC 1777 "The Lightweight Directory Access Protocol." This section gives an overview of LDAP from a user's perspective.

*What kind of information can be stored in the directory?* The LDAP directory service model is based on *entries*. An entry is a collection of *attributes* that has a name, called a *distinguished name* (DN). The DN is used to refer to the entry unambiguously. Each of the entry's attributes has a *type* and one or more *values*. The types are typically mnemonic strings, like "cn" for common name, or "mail" for email address. The values depend on what type of attribute it is. For example, a mail attribute might contain the value "babs@umich.edu". A jpegPhoto attribute would contain a photograph in binary JPEG/JFIF format.

*How is the information arranged?* In LDAP, directory entries are arranged in a hierarchical tree-like structure that reflects political, geographic and/or

organizational boundaries. Entries representing countries appear at the top of the tree. Below them are entries representing states or national organizations. Below them might be entries representing people, organizational units, printers, documents, or just about anything else you can think of. Figure 1 shows an example LDAP directory tree, which should help make things clear.

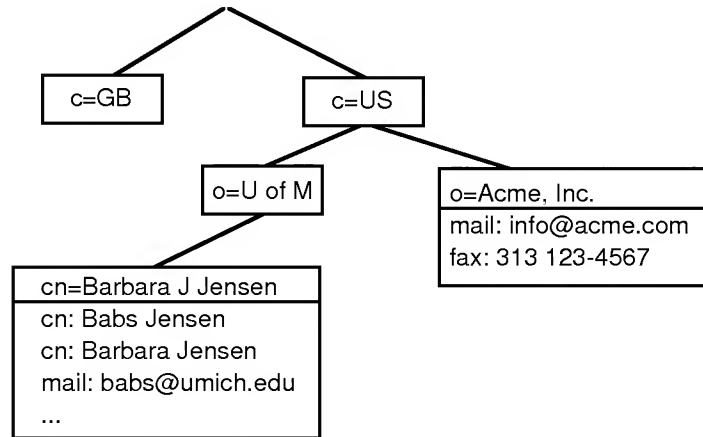


Figure 1: An example LDAP directory tree.

In addition, LDAP allows you to control which attributes are required and allowed in an entry through the use of a special attribute called `objectclass`. The values of the `objectclass` attribute determine the *schema rules* the entry must obey.

*How is the information referenced?* An entry is referenced by its distinguished name, which is constructed by taking the name of the entry itself (called the relative distinguished name, or RDN) and concatenating the names of its ancestor entries. For example, the entry for Barbara Jensen in the example above has an RDN of "cn=Barbara J Jensen" and a DN of "cn=Barbara J Jensen, o=U of M, c=US". The full DN format is described in RFC 1779, "A String Representation of Distinguished Names."

*How is the information accessed?* LDAP defines operations for interrogating and updating the directory. Operations are provided for adding and deleting an entry from the directory, changing an existing entry, and changing the name of an entry. Most of the time, though, LDAP is used to search for information in the directory. The LDAP search operation allows some portion of the directory to be searched for entries that match some criteria specified by a search filter. Information can be requested from each entry that matches the criteria.

For example, you might want to search the entire directory subtree below the University of Michigan for people with the name Barbara Jensen, retrieving the email address of each entry found. LDAP lets you do this easily. Or you might want to search the entries directly below the `c=US` entry for organizations with the string "Acme" in their name, and that have a fax number. LDAP lets you do this too. The next section describes in more detail what you can do with LDAP and how it might be useful to you.

*How is the information protected from unauthorized access?* Some directory services provide no protection, allowing anyone to see the information. LDAP provides a method for a client to authenticate, or prove its identity to a directory



server, paving the way for rich access control to protect the information the server contains.

### 1.3 How does LDAP work?

LDAP directory service is based on a *client-server* model. One or more LDAP servers contain the data making up the LDAP directory tree. An LDAP client connects to an LDAP server and asks it a question. The server responds with the answer, or with a pointer to where the client can get more information (typically, another LDAP server). No matter which LDAP server a client connects to, it sees the same view of the directory; a name presented to one LDAP server references the same entry it would at another LDAP server. This is an important feature of a global directory service, like LDAP.

### 1.4 What is slapd and what can it do?

*Slapd* is an LDAP directory server that runs on many different UNIX platforms. You can use it to provide a directory service of your very own. Your directory can contain pretty much anything you want to put in it. You can connect it to the global LDAP directory service, or run a service all by yourself. Some of *slapd*'s more interesting features and capabilities include:

**Choice of databases:** *Slapd* comes with three different backend databases you can choose from. They are LDBM, a high-performance disk-based database; SHELL, a database interface to arbitrary UNIX commands or shell scripts; and PASSWD, a simple password file database.

**Multiple database instances:** *Slapd* can be configured to serve multiple databases at the same time. This means that a single *slapd* server can respond to requests for many logically different portions of the LDAP tree, using the same or different backend databases.

**Generic database API:** If you require even more customization, *slapd* lets you write your own backend database easily. *Slapd* consists of two distinct parts: a front end that handles protocol communication with LDAP clients; and a backend that handles database operations. Because these two pieces communicate via a well-defined C API, you can write your own customized database backend to *slapd*.

**Access control:** *Slapd* provides a rich and powerful access control facility, allowing you to control access to the information in your database(s). You can control access to entries based on LDAP authentication information, IP address, domain name and other criteria.

**Threads:** *Slapd* is threaded for high performance. A single multi-threaded *slapd* process handles all incoming requests, reducing the amount of system overhead required. *Slapd* will automatically select the best thread support for your platform.

**Replication:** *Slapd* can be configured to maintain replica copies of its database. This master/slave replication scheme is vital in high-volume environments where a single *slapd* just doesn't provide the necessary availability or reliability.

**Configuration:** *Slapd* is highly configurable through a single configuration file which allows you to change just about everything you'd ever want to change. Configuration options have reasonable defaults, making your job much easier.

*Slapd* also has its limitations, of course. It does not currently handle aliases, which are part of the LDAP model. The main LDBM database backend does not handle range queries or negation queries very well. These features and more will be coming in a future release.

### **1.5 What about X.500?**

LDAP was originally developed as a front end to X.500, the OSI directory service. X.500 defines the Directory Access Protocol (DAP) for clients to use when contacting directory servers. DAP is a heavyweight protocol that runs over a full OSI stack and requires a significant amount of computing resources to run. LDAP runs directly over TCP and provides most of the functionality of DAP at a much lower cost.

This use of LDAP makes it easy to access the X.500 directory, but still requires a full X.500 service to make data available to the many LDAP clients being developed. As with full X.500 DAP clients, a full X.500 server is no small piece of software to run.

The stand-alone LDAP daemon, or *slapd*, is meant to remove much of the burden from the server side just as LDAP itself removed much of the burden from clients. If you are already running an X.500 service and you want to continue to do so, you can probably stop reading this guide, which is all about running LDAP via *slapd*, without running X.500. If you are not running X.500, want to stop running X.500, or have no immediate plans to run X.500, read on.

It is possible to replicate data from a *slapd* directory server to an X.500 DSA, which allows your organization to make your data available as part of the global X.500 directory service on a "read-only" basis. This is discussed in section 11.6.

Another way to make data in a *slapd* server available to the X.500 community would be by using a X.500 DAP to LDAP gateway. At this time, no such software has been written (to the best of our knowledge), but hopefully some group will see fit to write such a gateway.

### **1.6 What is slurpd and what can it do?**

*Slurpd* is a UNIX daemon that helps *slapd* provide replicated service. It is responsible for distributing changes made to the master *slapd* database out to the various *slapd* replicas. It frees *slapd* from having to worry that some replicas might be down or unreachable when a change comes through; *slurpd* handles retrying failed requests automatically. *Slapd* and *slurpd* communicate through a simple text file that is used to log changes.